Network Working Group
Request for Comments: 4395
Obsoletes: 2717, 2718
BCP: 115
Category: Best Current Practice

T. Hansen
AT&T Laboratories
T. Hardie
Qualcomm, Inc.
L. Masinter
Adobe Systems

February 2006

Guidelines and Registration Procedures for New URI Schemes

Status of This Memo

Copyright Notice

Abstract

This document provides guidelines and recommendations for the definition of Uniform Resource Identifier (URI) schemes. It also updates the process and IANA registry for URI schemes. It obsoletes both RFC 2717 and RFC 2718.

Table of Contents

1. Introduction

The Uniform Resource Identifier (URI) protocol element and generic syntax is defined by RFC 3986 [5]. Each URI begins with a scheme name, as defined by Section 3.1 of RFC 3986, that refers to a specification for identifiers within that scheme. The URI syntax provides a federated and extensible naming system, where each scheme's specification may further restrict the syntax and semantics of identifiers using that scheme. This document provides guidelines for the definition of new URI schemes, for consideration by those who are defining, registering, or evaluating those definitions, as well as a process and mechanism for registering URI schemes within the IANA URI scheme registry. The registry has two parts: 'provisional' and 'permanent', with different requirements. Guidelines and requirements for both parts are given.

This document obsoletes both RFCs 2717 [7] and 2718 [8]. RFCs 2717 and 2718 drew a distinction between 'locators' (identifiers used for accessing resources available on the Internet) and 'names' (identifiers used for naming possibly abstract resources, independent of any mechanism for accessing them). The intent was to use the designation "URL" (Uniform Resource Locator) for those identifiers that were locators and "URN" (Uniform Resource Name) for those identifiers that were names. In practice, the line between 'locator' and 'name' has been difficult to draw: locators can be used as names, and names can be used as locators.

As a result, recent documents have used the term "URI" for all resource identifiers, avoiding the term "URL" and reserving the term "URN" explicitly for those URIs using the "urn" scheme name (RFC 2141 [2]). URN "namespaces" (RFC 3406 [9]) are specific to the "urn" scheme and not covered explicitly by this document.

RFC 2717 defined a set of registration trees in which URI schemes could be registered, one of which was called the IETF Tree, to be managed by IANA. RFC 2717 proposed that additional registration trees might be approved by the IESG. However, no such registration trees have been approved.

This document eliminates RFC 2717's distinction between different 'trees' for URI schemes; instead there is a single namespace for registered values. Within that namespace, there are values that are approved as meeting a set of criteria for URI schemes. Other scheme names may also be registered provisionally, without necessarily meeting those criteria. The intent of the registry is to:

- o provide a central point of discovery for established URI scheme names, and easy location of their defining documents;
- o discourage use of the same URI scheme name for different purposes;
- o help those proposing new URI scheme names to discern established trends and conventions, and avoid names that might be confused with existing ones;
- o encourage registration by setting a low barrier for provisional registrations.

RFC 3987 [6] introduced a new protocol element, the Internationalized Resource Identifier (IRI), and defined a mapping between URIs and IRIs. There is no separate registry or registration process for IRIs. Those who wish to describe resource identifiers that are useful as IRIs should define the corresponding URI syntax, and note that the IRI usage follows the rules and transformations defined in RFC 3987.

Within this document, the key words MUST, MAY, SHOULD, REQUIRED, RECOMMENDED, and so forth are used within the general meanings established in RFC 2119 [1], within the context that they are requirements on future registration documents.

2. Guidelines for Permanent URI Scheme Definitions

This section gives considerations for new URI schemes. Meeting these guidelines is REQUIRED for permanent URI scheme registration. Meeting these guidelines is also RECOMMENDED for provisional registration, as described in Section 3.

2.1. Demonstratable, New, Long-Lived Utility

The use and deployment of new URI schemes in the Internet infrastructure is costly; some parts of URI processing may be scheme-dependent, and deployed software already processes URIs of well-known schemes. Introducing a new URI scheme may require additional software, not only for client software and user agents but also in additional parts of the network infrastructure (gateways, proxies, caches) [11]. URI schemes constitute a single, global namespace; it is desirable to avoid contention over use of short, mnemonic scheme names. For these reasons, the unbounded registration of new schemes is harmful. New URI schemes SHOULD have clear utility to the broad Internet community, beyond that available with already registered URI schemes.

2.2. Syntactic Compatibility

RFC 3986 [5] defines the generic syntax for all URI schemes, along with the syntax of common URI components that are used by many URI schemes to define hierarchical identifiers. All URI scheme specifications MUST define their own syntax such that all strings matching their scheme-specific syntax will also match the <absolute-URI> grammar described in Section 4.3 of RFC 3986.

New URI schemes SHOULD reuse the common URI components of RFC 3986 for the definition of hierarchical naming schemes. However, if there is a strong reason for a URI scheme not to use the hierarchical syntax, then the new scheme definition SHOULD follow the syntax of previously registered schemes.

URI schemes that are not intended for use with relative URIs SHOULD avoid use of the forward slash "/" character, which is used for hierarchical delimiters, and the complete path segments "." and ".." (dot-segments).

Avoid improper use of "//". The use of double slashes in the first part of a URI is not an artistic indicator that what follows is a URI: Double slashes are used ONLY when the syntax of the URI's <scheme-specific-part> contains a hierarchical structure as described in RFC 3986. In URIs from such schemes, the use of double slashes indicates that what follows is the top hierarchical element for a naming authority. (See Section 3.2 of RFC 3986 for more details.) URI schemes that do not contain a conformant hierarchical structure in their <scheme-specific-part> SHOULD NOT use double slashes following the "<scheme>:" string.

New URI schemes SHOULD clearly define the role of RFC 3986 [5] reserved characters in URIs of the scheme being defined. The syntax of the new scheme should be clear about which of the "reserved" set of characters (as defined in RFC 3986) are used as delimiters within the URIs of the new scheme, and when those characters must be escaped, versus when they may be used without escaping.

2.3. Well-Defined

While URIs may or may not be useful as locators in practice, a URI scheme definition itself MUST be clear as to how it is expected to function. Schemes that are not intended to be used as locators SHOULD describe how the resource identified can be determined or accessed by software that obtains a URI of that scheme.

For schemes that function as locators, it is important that the mechanism of resource location be clearly defined. This might mean different things depending on the nature of the URI scheme.

In many cases, new URI schemes are defined as ways to translate between other namespaces or protocols and the general framework of URIs. For example, the "ftp" URI scheme translates into the FTP protocol, while the "mid" URI scheme translates into a Message-ID identifier of an email message. For such schemes, the description of the mapping must be complete, and in sufficient detail so that the mapping in both directions is clear: how to map from a URI into an identifier or set of protocol actions or name in the target namespace, and how legal values in the base namespace, or legal protocol interactions, might be represented in a valid URI. In particular, the mapping should describe the mechanisms for encoding binary or character strings within valid character sequences in a URI (See Section 2.6 for guidelines). If not all legal values or protocol interactions of the base standard can be represented using the URI scheme, the definition should be clear about which subset are allowed, and why.

2.4. Definition of Operations

As part of the definition of how a URI identifies a resource, a URI scheme definition SHOULD define the applicable set of operations that may be performed on a resource using the URI as its identifier. A model for this is HTTP; an HTTP resource can be operated on by GET, POST, PUT, and a number of other operations available through the HTTP protocol. The URI scheme definition should describe all well-defined operations on the URI identifier, and what they are supposed to do.

Some URI schemes don’t fit into the "information access" paradigm of URIs. For example, "telnet" provides location information for initiating a bi-directional data stream to a remote host; the only operation defined is to initiate the connection. In any case, the operations appropriate for a URI scheme should be documented.

Note: It is perfectly valid to say that "no operation apart from GET is defined for this URI". It is also valid to say that "there’s only one operation defined for this URI, and it’s not very GET-like". The important point is that what is defined on this scheme is described.

2.5. Context of Use

In general, URIs are used within a broad range of protocols and applications. Most commonly, URIs are used as references to resources within directories or hypertext documents, as hyperlinks to

other resources. In some cases, a URI scheme is intended for use within a different, specific set of protocols or applications. If so, the scheme definition SHOULD describe the intended use and include references to documentation that define the applications and/or protocols cited.

2.6. Internationalization and Character Encoding

When describing URI schemes in which (some of) the elements of the URI are actually representations of human-readable text, care should be taken not to introduce unnecessary variety in the ways in which characters are encoded into octets and then into URI characters; see RFC 3987 [6] and Section 2.5 of RFC 3986 [5] for guidelines. If URIs of a scheme contain any text fields, the scheme definition MUST describe the ways in which characters are encoded, and any compatibility issues with IRIs of the scheme.

2.7. Clear Security Considerations

Definitions of URI schemes MUST be accompanied by a clear analysis of the security implications for systems that use the URI scheme; this follows the practice of Security Consideration sections within IANA registrations [3].

In particular, Section 7 of RFC 3986 [5] describes general security considerations for URI schemes. The definition of an individual URI scheme should note which of these apply to the specified scheme.

2.8. Scheme Name Considerations

Section 3.1 of RFC 3986 defines the syntax of a URI scheme name. New scheme registrations MUST comply. Note that although scheme names are case insensitive, scheme names MUST be registered using lowercase letters.

URI scheme names should be short, but also sufficiently descriptive and distinguished to avoid problems.

Avoid names or other symbols that might cause problems with rights to use the name in IETF specifications and Internet protocols. For example, be careful with trademark and service mark names. (See Section 7.4 of RFC 3978 [4].)

Avoid using names that are either very general purpose or associated in the community with some other application or protocol. Avoid scheme names that are overly general or grandiose in scope (e.g., that allude to their "universal" or "standard" nature when the described namespace is not.)

Organizations that desire a private name space for URI scheme names are encouraged to use a prefix based on their domain name, expressed in reverse order. For example, a URI scheme name of com-example-info might be registered by the vendor that owns the example.com domain name.

3. Guidelines for Provisional URI Scheme Registration

While the guidelines in Section 2 are REQUIRED for permanent registration, they are RECOMMENDED for provisional registration. For a provisional registration, the following are REQUIRED:

- The scheme name meets the syntactic requirements of Section 2.8.
- There is not already an entry with the same URI scheme name. (In the unfortunate case that there are multiple, different uses of the same scheme name, the IESG may approve a request to modify an existing entry to note the separate use.)
- Contact information identifying the person supplying the registration is included. Previously unregistered URI schemes discovered in use may be registered by third parties on behalf of those who created the URI scheme; in this case, both the registering party and the scheme creator SHOULD be identified.
- If no permanent, citable specification for the URI scheme definition is included, credible reasons for not providing it should be given.
- A valid Security Considerations section, as required by Section 6 of [3].
- If the scheme definition does not meet the guidelines laid out in Section 2, the differences and reasons SHOULD be noted.

4. Guidelines for Historical URI Scheme Registration

In some circumstances, it is appropriate to note a URI scheme that was once in use or registered but for whatever reason is no longer in common use or the use is not recommended. In this case, it is possible for an individual to request that the URI scheme be registered (newly, or as an update to an existing registration) as 'historical'. Any scheme that is no longer in common use MAY be designated as historical; the registration should contain some indication to where the scheme was previously defined or documented.

5. URI Scheme Registration Procedure

5.1. General

The URI registration process is described in the terminology of [3]. The registration process is an optional mailing list review, followed by "Expert Review". The registration request should note the desired status. The Designated Expert will evaluate the request against the criteria of the requested status. In the case of a permanent registration request, the Designated Expert may:

- Accept the URI scheme name for permanent registration.
- Suggest provisional registration instead.
- Request IETF review and IESG approval; in the meanwhile, suggest provisional registration.

URI scheme definitions contained within other IETF documents (Informational, Experimental, or Standards-Track RFCs) must also undergo Expert Review; in the case of Standards-Track documents, permanent registration status approval is required.

5.2. Registration Procedures

Someone wishing to register a URI scheme SHOULD:

1. Check the IANA URI scheme registry to see whether or not there is already an entry for the desired name. If there is already an entry under the name, choose a different URI scheme name.
2. Prepare a URI scheme registration template, as specified in Section 5.4. The URI scheme registration template may be contained in an Internet Draft, alone or as part of some other protocol specification. The template may also be submitted in some other form (as part of another document or as a stand-alone document), but the contents will be treated as an "IETF Contribution" under the guidelines of RFC 3978 [4].
3. Send a copy of the template or a pointer to the containing document (with specific reference to the section with the template) to the mailing list uri-review@ietf.org, requesting review. In addition, request review on other mailing lists as appropriate. For example, general discussion of URI syntactical issues could be discussed on uri@w3.org; schemes for a network protocol could be discussed on a mailing list for that protocol. Allow a reasonable time for discussion and comments. Four weeks is reasonable for a permanent registration requests.
4. Respond to review comments and make revisions to the proposed registration as needed to bring it into line with the guidelines given in this document.

5. Submit the (possibly updated) registration template (or pointer to document containing it) to IANA at iana@iana.org, specifying whether 'permanent' or 'provisional' registration is requested.

Upon receipt of a URI scheme registration request,

1. IANA checks the submission for completeness; if sections are missing or citations are not correct, IANA rejects the registration request.
2. IANA checks the current registry for a entry with the same name; if such a registry exists, IANA rejects the registration request.
3. IANA requests Expert Review of the registration request against the corresponding guidelines.
4. The Designated Expert may request additional review or discussion, as necessary.
5. If Expert Review recommends registration 'provisional' or 'permanent' registration, IANA adds the registration to the appropriate registry.
6. Unless Expert Review has explicitly rejected the registration request within two weeks, IANA should automatically add the registration in the 'provisional' registry.

Either based on an explicit request or independently initiated, the Designated Expert or IESG may request the upgrade of a 'provisional' registration to a 'permanent' one. In such cases, IANA should move the corresponding entry from the provisional registry.

5.3. Change Control

Registrations may be updated in each registry by the same mechanism as required for an initial registration. In cases where the original definition of the scheme is contained in an IESG-approved document, update of the specification also requires IESG approval.

Provisional registrations may be updated by the original registrant or anyone designated by the original registrant. In addition, the IESG may reassign responsibility for a provisional registration scheme, or may request specific changes to a scheme registration. This will enable changes to be made to schemes where the original registrant is out of contact, or unwilling or unable to make changes.

Transition from 'provisional' to 'permanent' status may be requested and approved in the same manner as a new 'permanent' registration. Transition from 'permanent' to 'historical' status requires IESG approval. Transition from 'provisional' to 'historical' may be requested by anyone authorized to update the provisional registration.

5.4. URI Scheme Registration Template

This template describes the fields that must be supplied in a URI scheme registration request:

URI scheme name.
: See Section 2.8 for guidelines.

Status.
: This reflects the status requested, and should be one of 'permanent', 'provisional', or 'historical'.

URI scheme syntax.
: See Section 2.2 for guidelines.

URI scheme semantics.
: See Section 2.3 and Section 2.4 for guidelines.

Encoding considerations.
: See Section 2.3 and Section 2.6 for guidelines.

Applications/protocols that use this URI scheme name.
: Applications and/or protocols that use this URI scheme name; see Section 2.5.

Interoperability considerations.
: If you are aware of any details regarding your scheme that might impact interoperability, please identify them here. For example: proprietary or uncommon encoding method; inability to support multibyte character sets; incompatibility with types or versions of any underlying protocol.

Security considerations.
: See Section 2.7 for guidelines.

Contact.
: Person (including contact information) to contact for further information.

Author/Change controller.
: Person (including contact information) authorized to change this, if a provisional registration.

References.
: Include full citations for all referenced documents. Registration templates for provisional registration may be included in an Internet Draft; when the documents expire or are approved for publication as an RFC, the registration will be updated.

6. IANA Considerations

This document replaces the current "URL Scheme" registry with a new Uniform Resource Identifier scheme registry, and establishes a new registration template and a new process for registration. The process is based on [3] "Expert Review" with an initial (optional) mailing list review.

The template has an additional field for the status of the URI name scheme, and the procedures for entering new name schemes have been augmented. Section 5 establishes the process for new URI scheme registration.

To transition to the new registry, all URL name schemes in the existing table should be entered as URI schemes, with 'permanent' status.

7. Security Considerations

All registered values are expected to contain accurate security consideration sections; 'permanent' registered scheme names are expected to contain complete definitions.

Information concerning possible security vulnerabilities of a protocol may change over time. Consequently, claims as to the security properties of a registered URI scheme may change as well. As new vulnerabilities are discovered, information about such vulnerabilities may need to be attached to existing documentation, so that users are not misled as to the true security properties of a registered URI scheme.

8. Acknowledgements

Many thanks to Paul Hoffmann, Ira McDonald, Roy Fielding, Stu Weibel, Tony Hammond, Charles Lindsey, Mark Baker, and other members of the uri@w3.org mailing list for their comments on earlier versions.

Parts of this document are based on [7], [8] and [10]. Some of the ideas about use of URIs were taken from the "Architecture of the World Wide Web" [11].

9. References

9.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Moats, R., "URN Syntax", RFC 2141, May 1997.

[3] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[4] Bradner, S., "IETF Rights in Contributions", BCP 78, RFC 3978, March 2005.

[5] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

[6] Duerst, M. and M. Suignard, "Internationalized Resource Identifiers (IRIs)", RFC 3987, January 2005.

9.2. Informative References

[7] Petke, R. and I. King, "Registration Procedures for URL Scheme Names", BCP 35, RFC 2717, November 1999.

[8] Masinter, L., Alvestrand, H., Zigmond, D., and R. Petke, "Guidelines for new URL Schemes", RFC 2718, November 1999.

[9] Daigle, L., van Gulik, D., Iannella, R., and P. Faltstrom, "Uniform Resource Names (URN) Namespace Definition Mechanisms", BCP 66, RFC 3406, October 2002.

[10] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", BCP 90, RFC 3864, September 2004.

[11] W3C Technical Architecture Group, "Architecture of the World Wide Web, Volume One", December 2004, <http://www.w3.org/TR/webarch/>.

Authors' Addresses

Tony Hansen
AT&T Laboratories
200 Laurel Ave.
Middletown, NJ 07748
USA

EMail: tony+urireg@maillennium.att.com

Ted Hardie
Qualcomm, Inc.
675 Campbell Technology Parkway
Campbell, CA
USA

EMail: hardie@qualcomm.com

Larry Masinter
Adobe Systems
345 Park Ave
San Jose, CA 95110
US

Phone: +1 408 536 3024
EMail: LMM@acm.org
URI: http://larry.masinter.net

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).